AFRL-RY-WP-TR-2012-0393

FRAMEWORK FOR EFFICIENT ALGORITHMS IN PLANAR NETWORKS AND BEYOND

Erik D. Demaine
MIT

Mohammad T. Hajiaghayi
University of Maryland

DECEMBER 2012
Final Report

Approved for public release; distribution unlimited

AIR FORCE RESEARCH LABORATORY
SENSORS DIRECTORATE
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7320
AIR FORCE MATERIEL COMMAND
UNITED STATES AIR FORCE

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nationals.

Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RY-WP-TR-2012-0393 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

*//Signature//

MICHAEL BAKICH, Program Manager
Layered Sensing Performance Characterization Branch

//Signature//

CHRISTINA SCHUTTE, Branch Chief
Layered Sensing Performance Characterization Branch

//Signature//

DOUG HAGER, Deputy
Layered Sensing Exploitation Division
Sensors Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

*Disseminated copies will show "//Signature//" stamped or typed above the signature blocks.

REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE ***(DD-MM-YY)***	**2. REPORT TYPE**	**3. DATES COVERED** ***(From - To)***
December 2012	Final	15 September 2011 – 15 September 2012

4. TITLE AND SUBTITLE

FRAMEWORK FOR EFFICIENT ALGORITHMS IN PLANAR NETWORKS AND BEYOND

5a. CONTRACT NUMBER FA8650-11-1-7162

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER 61101E

6. AUTHOR(S)

Erik D. Demaine (MIT)
Mohammad T. Hajiaghayi (University of Maryland)

5d. PROJECT NUMBER 1000

5e. TASK NUMBER 11

5f. WORK UNIT NUMBER Y0ME

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

MIT
77 Massachusetts Avenue
Cambridge, MA 02139-4307

University of Maryland

8. PERFORMING ORGANIZATION REPORT NUMBER

AFRL-RY-WP-TR-2012-0393

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Research Laboratory
Sensors Directorate
Wright-Patterson Air Force Base, OH 45433-7320
Air Force Materiel Command
United States Air Force

Defense Advanced Research Projects Agency (DARPA/DSO)
675 North Randolph Street
Arlington, VA 22203-2114

10. SPONSORING/MONITORING AGENCY ACRONYM(S)

AFRL/RYAA

11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)

AFRL-RY-WP-TR-2012-0393

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

13. SUPPLEMENTARY NOTES

This material is based on research sponsored by Air Force Research Laboratory (AFRL) and the Defense Advanced Research Agency (DARPa) under agreement number FA8650-11-1-7162. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes not withstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFRL and DARPA or the U.S. Government. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nationals. Report contains color.

14. ABSTRACT

During the one year of this grant, we proved several breakthrough results to develop frameworks for approximation and fixed-parameter algorithms in planar and nearly planar graphs. In this final report, we detal these results and summarize the supported efforts of students and visitors.

15. SUBJECT TERMS

Algorithms, graphs, hard problems

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT:**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON (Monitor)**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Michael Bakich
Unclassified	Unclassified	Unclassified	SAR	10	**19b. TELEPHONE NUMBER** ***(Include Area Code)*** N/A

Progress Report for DARPA/AFRL grant FA8650-11-1-7162 "Frameworks for Efficient Algorithms in Planar Networks and Beyond"

Erik D. Demaine Mohammad T. Hajiaghayi

Final Report: Sept. 2011–Aug. 2012

During the one year of this grant, we proved several breakthrough results to develop frameworks for approximation and fixed-parameter algorithms in planar and nearly planar graphs. In this final report, we detail these results and summarize the supported efforts of students and visitors.

Contraction Decomposition

We solved Research Challenge 4 described in the proposal, and published it in the premiere conference in theoretical computer science:

> Erik D. Demaine, MohammadTaghi Hajiaghayi, and Ken-ichi Kawarabayashi, "Contraction decomposition in H-minor-free graphs and algorithmic applications", in *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC 2011)*, 2011, pages 441–450. http://erikdemaine.org/papers/ContractionMinorFree_STOC2011/

High-level description of result. Our goal in this research is to develop frameworks for solving NP-hard problems, very close to optimality. (Solving NP-hard problems to optimality is broadly believed to be impossible, so this is the best we can hope for.) The domain of algorithmic problems we look at involve *graphs*.

In general, a graph represents pairwise relationships among arbitrary objects called *nodes*. For example, the nodes might represent people and the edges might represent relationships between two people; or the nodes might represent computers and the edges might represent network links between computers (either physical wires or the ability to communicate wirelessly); or the nodes might represent cities and the edges might represent roads connecting cities.

A *planar* graph is a graph that can be drawn in two dimensions without any of the edges crossing each other. Many graphs of practical interest are planar, or nearly planar. For example, computer networks and road networks are "drawn" on Earth, and thus generally avoid crossings.

However, real networks often have some crossings. For example, road networks introduce crossings using overpasses; and computer networks can introduce crossings by digging deeper underground. Therefore we would like an algorithmic theory for solving problems on "nearly planar" graphs.

The strongest known form of "almost planar" is called K_k-*minor-free*. A graph is K_k-minor-free if it has k connected subgraphs, each representing a "supernode", which do not overlap each other, such that the graph has an edge connecting every pair of supernodes. Planar graphs are K_5-minor-free. By allowing k to grow, we allow many crossings, provided they are "well-behaved".

At high level, our new result, called *contraction decomposition*, decomposes any K_k-minor-free graph into "algorithmically simple" structures, which allows us to solve many algorithmic problems very close to optimality.

Specifically, the notion of "algorithmically simple" graphs we use is *bounded-treewidth*. Most NP-hard problems become easy when considering just *tree* graphs, that is, connected graphs with no cycles or loops. A bounded-treewidth graph is a graph that is similar to a tree, just a little "thicker". Most NP-hard problems also become easy on bounded-treewidth graphs.

Our decomposition is based on a graph operation called *contraction*. Contraction brings together the two ends of an edge, combining these two nodes into a single node. Many problems become easier after performing a contraction. For example, in the venerable Traveling Salesman Problem,[1] the goal is to find a minimum-length tour that visits every node in the graph. If we contract an edge in the graph, the tour can essentially visit two nodes for the price of one, and all previous tours still "work", and thus the optimal solution can only get better.

Now can precisely specify how our contraction decomposition result works. Given a K_k-minor-free graph and a positive integer p, an algorithm partitions the edges into p groups such that contracting all the edges in any one group (and keeping the other $p-1$ groups) results in a bounded-treewidth graph. Thus, there are p different graphs we can get, by contracting any one of the groups, and every one of them is a bounded-treewidth graph, and thus our problem becomes easy to solve algorithmically.

This result leads to a very general approach to finding solutions that are very close to optimal. Consider any problem where the goal is to find a collection of edges with various constraints (for example, forming a tour in the Traveling Salesman Problem). Any one of the p groups will, on average, have a $1/p$ fraction of the edges in the optimal solution. In particular, some one of the p groups will have at most a $1/p$ fraction of the edges of the optimal solution. An algorithm does not know which group this is, but it can try each of the p, solve the problem perfectly on the resulting bounded-treewidth graph, and then take the best solution so found.

This approach generally leads to a solution within a factor of $1 + 1/p$ of the optimal solution. So, for example, if we want to have at most 1% error, we simply set $p = 100$. In general, we can set the error to be as small as we want, though the running time will increase as the error decreases.

This contraction-decomposition result represents the culmination of several papers over several years, related to a weaker form of decomposition called "deletion decomposition" as well as similar contraction decomposition results but for less general graphs.

Multiway Cut

We solved Research Challenge 5 described in the proposal, about the multiway cut problem, and published it in the premiere conference in algorithms:

MohammadHossein Bateni, MohammadTaghi Hajiaghayi, Philip N. Klein, Claire Math-

[1]See, e.g., the recent *New York Times* article: http://campaignstops.blogs.nytimes.com/2011/12/21/the-problem-of-the-traveling-politician/

ieu, “A polynomial-time approximation scheme for planar multiway cut”, in *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, 2012, pages 639–655.

High-level description of result. Given an undirected graph with edge lengths and a subset of nodes (called the *terminals*), the *multiway cut* (also called the *multi-terminal cut*) problem asks for a subset of edges, with minimum total length, whose removal disconnects each terminal from all others.

Multiway cut is an old problem: a variant of multiway cut was first proposed in T. C. Hu's 1969 book. It is a natural problem: it generalizes the problem of finding a minimum-length *st*-cut. It is an important problem: applications proposed for multiway cut include image processing, chip design as well as parallel and distributed computing. It is a well-studied problem: the study of its computational complexity was inaugurated in 1983 by Dahlhaus, Johnson, Papadimitriou, Seymour, and Yannakakis (the journal version was published in 1994). Their results, already highlighting the case of planar graphs because of its relevance to image processing, have guided the agenda for subsequent research. In particular for planar graphs, they solve the multiway cut problem in polynomial time for fixed k and show the problem is NP-Hard when k is unbounded. The main remaining open problem in this work is to obtain the best approximation factor for multiway cut on planar graphs (they show for general graphs, there is a simple 2-approximation).

In this breakthrough result, we resolve the important open problem above after 29 years by presenting a Polynomial-Time Approximation Scheme (PTAS)[2] for *planar multiway cut.* We prove the result by building a novel “spanner” for multiway cut on planar graphs which is of independent interest.

More precisely the overview of the paradigm is as follows. First we start by constructing a *skeleton*: roughly speaking, a subgraph of length at most a constant times the optimal cost, and such that each connected component of a $(1+\epsilon)$-approximate solution intersects exactly one connected component of the spanner. Given the skeleton, we can then construct a *spanner*: a subgraph of length $g(\epsilon)$ times the optimal cost that contains a $(1+\epsilon)$-approximate solution. Given the spanner, in planar graphs there is a reduction to *instances of bounded treewidth*, on which the problem, in most cases, is easily solvable by dynamic programming. This paradigm has been used recently to obtain PTASs for other problems like Steiner tree and Steiner forest on planar graphs.

For the multiway cut problem, we work in the dual and reuse suitably adapted versions of the general paradigm above however still we needed to overcome several substantial difficulties. Below we mention at least three additional ideas needed to define and build a multiway cut skeleton.

First, the optimal solution may be *not connected* in the dual. Moreover, the subset of edges “serving” a terminal need not be connected; it may consist of several disjoint cycles. Given a terminal t, we focus on the dual cycle C^* that, in the dual of the optimal solution, minimally encloses t. We include in the skeleton a small number of cycles enclosing t to guarantee that at least one, $C(t)$, intersects C^*.

Second, we add “ears” and precede the construction with some preprocessing to guarantee that, in the connected component of $C(t)$, when we consider the face containing t, we are able to “guess” the identity of t from only a small number of possible candidates.

Third, in the optimal solution, each edge serves to disconnect one particular pair of terminals. The set of edges in the solution determines a subset of disconnection constraints, where in the dual

[2] $1+\epsilon$-approximation algorithm for arbitrary constant ϵ.

optimal solution, each terminal t is minimally enclosed by a cycle whose purpose is to disconnect t from some specific other terminals. Very roughly speaking, the algorithm guesses those disconnection constraints and then constructs a minimal solution that is guaranteed to satisfy them. Unfortunately this does not imply global feasibility in general. We deal with this problem using the technique of prize-collecting clustering. This is a new use of the prize-collecting clustering theorem.

Coordinated Movement Problems

We made major progress on Research Challenge 7 described in the proposal, about coordinated movement problems, and published it in the APPROX conference:

> Piotr Berman, Erik D. Demaine, and Morteza Zadimoghaddam, "$O(1)$-Approximations for Maximum Movement Problems", in *Proceedings of the 14th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, pages 62–74.

High-level motivation. A central problem in swarm robotics is to reconfigure the robots into an arrangement with a desired property. For example, in the *connectivity* goal, the proximity of the robots should form a connected graph. Two motivations for this goal are forming a connected data network with short-range radios, and forming a connected physical network for transporting materials. In the first situation, the robots initially communicate via another channel, e.g., via slow and/or power-intensive long-distance communication (such as satellite or the same radios with power turned up high), or via two traversals by aircraft to locate robots and disseminate locations. Another connectivity goal is *Steiner connectivity*, where a subset of the robots should form a connected network that contains k stationary nodes (such as buildings or sensors) which need to be interconnected. In both of these problems, we suppose that we know the initial locations of robots, and that we have a map of the environment the robots can traverse and defining proximity among the robots. Our goal is to move the robots as quickly as possible into a configuration with the desired property, i.e., to minimize the maximum motion required by any robot.

New results. We study approximation of the fastest way to bring robots into a connected configuration. The best previous approximation results (in a previous paper of ours) left a gap in the best possible approximation ratio between $2 - \epsilon$ and $O(\sqrt{n})$. In our new work, we close this gap up to constant factors, by obtaining the first constant-factor approximation algorithm.

An ingredient in this result is a constant-factor approximation algorithm for the *s-t path movement problem*, where the goal is for the robots to form a connected path connecting two given locations s and t. This result is also a breakthrough, as again the best possible approximation ratio was previously only known to be between $2 - \epsilon$ and $O(\sqrt{n})$. We use our approximation algorithm for s-t path problem as a black box in our solution to the connectivity problem.

We also study the Steiner connectivity movement problem, which is a natural generalization of the s-t path movement problem. Here we are given a set T of terminal vertices, and the goal is to move some of our robots to interconnect all terminal vertices. For $|T| = O(\log n)$, we present an $O(|T|)$-approximation algorithm for the Steiner connectivity movement problem with the maximum movement objective, again using our approximation algorithm for s-t path. Note that we cannot hope to approximate the Steiner connectivity movement problem for arbitrary $|T|$, because even deciding whether there is a feasible solution with the available robots is the NP-hard node-weighted Steiner tree problem.

Our algorithms introduce several new techniques for approximating movement problems with the maximum movement objective, which we hope will extend to other movement problems. In general, the connectivity problems would be easy if we allowed approximating the number of robots (via resource augmentation) in addition to the cost. But robots cannot (yet) replicate themselves, so resource augmentation is not very useful. We develop powerful tools to resolve multiple desires for the location of a single robot.

Refined fixed-parameter algorithms for movement problems. We also completed our paper:

> Erik D. Demaine, MohammadTaghi Hajiaghayi, and Dániel Marx, "Minimizing Movement: Fixed-Parameter Tractability"

which we plan to submit to the journal *ACM Transactions on Algorithms.*

This paper considers exact solutions to the coordinated movement problems in the fixed-parameter context, where the number of moving robots is relatively small. Specifically, we develop general efficient algorithms for a broad family of movement problems. Furthermore, we show our results are tight by characterizing, in a very general setting, the line between fixed-parameter tractability and intractability. It turns out that the notion of treewidth plays an important role in defining this boundary line. Specifically we show that, for problems closed under edge addition (i.e., adding an edge to the connectivity graph cannot destroy a solution), the complexity of the problem depends solely on whether the edge-deletion minimal graphs of the property have bounded treewidth. If they all have bounded treewidth, we show how to solve a very general formulation of the problem with an efficient fixed-parameter algorithm. If they have unbounded treewidth, we show that even very simple questions are $W[1]$-hard, meaning there is no efficient fixed-parameter algorithm under the standard parameterized complexity assumption FPT $\neq$ W[1]. Thus we obtain a very precise characterization and general framework for fixed-parameter solutions to movement problems.

The new version of the paper completes, details, and extends our previous results in this setting. In particular, we developed a new simpler framework for describing our results, and developed especially efficient solutions for planar graphs.

General Approximation Schemes

We are currently preparing the journal version of our work on general techniques for approximation schemes in our bidimensionality framework, as described in Subtask 1 of the proposal:

> Erik D. Demaine and MohammadTaghi Hajiaghayi, "Bidimensionality: New Connections between FPT Algorithms and PTASs"

High-level description of results. We demonstrate a new connection between fixed-parameter tractability and approximation algorithms for combinatorial optimization problems on planar graphs and their generalizations. Specifically, we extend the theory of so-called "bidimensional" problems to show that essentially all such problems have both subexponential fixed-parameter algorithms and PTASs. Bidimensional problems include e.g. feedback vertex set, vertex cover, minimum maximal matching, face cover, a series of vertex-removal problems, dominating set, edge dominating set, r-dominating set, diameter, connected dominating set, connected edge dominating set, and connected r-dominating set. We obtain PTASs for all of these problems in planar graphs and certain generalizations; of particular interest are our results for the two well-known problems of connected dominating set and general feedback vertex set for planar graphs and their generalizations, for which PTASs were not known to exist. Our techniques generalize and in some sense unify the two main previous approaches for designing PTASs in planar graphs, namely, the Lipton-Tarjan separator approach [FOCS'77] and the Baker layerwise decomposition approach [FOCS'83]. In particular, we replace the notion of separators with a more powerful tool from the bidimensionality theory, enabling the first approach to apply to a much broader class of minimization problems than previously possible; and through the use of a structural backbone and thickening of layers we demonstrate how the second approach can be applied to problems with a "nonlocal" structure.

This work has formed the foundation of several follow-up papers:

> Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos, "(Meta) Kernelization", in *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, 2009, pages 629–638.
>
> Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos, "Bidimensionality and kernels", in *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, 2010, pages 503–510.
>
> Fedor V. Fomin, Daniel Lokshtanov, Venkatesh Raman, and Saket Saurabh, "Bidimensionality and EPTAS", in *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, 2011, pages 748–759.
>
> Fedor V. Fomin, Daniel Lokshtanov, and Saket Saurabh, "Bidimensionality and geometric graphs", in *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, 2012, pages 1563–1575.
>
> Philip N. Klein, "A linear-time approximation scheme for planar weighted TSP", in *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, 2005, pages 647–656.
>
> Philip N. Klein, "A subset spanner for planar graphs, with application to subset TSP", in *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, 2006, pages 749–756.
>
> Glencora Borradaile, Philip Klein, and Claire Mathieu, "An $O(n \log n)$ approximation scheme for Steiner tree in planar graphs", *ACM Transactions on Algorithms*, volume 5, number 3, July 2009.

New work. Our original paper on this topic omitted several details that we have been filling in. In particular, we have formally defined the separation property for contraction-bidimensional problems, which encompass many practical problems of interest.

Supported Students and Visitors

In the past year, this grant has supported research assistantships, travel, and/or equipment for the following PhD students: Rajesh Chitnis, Marek Cygan, Sarah Eisenstat, Reza Khani, and Morteza Zadimoghaddam. In addition, it supported the travel of the following visitors: Howard Karloff, Christian Sommer, Siamak Tazari, and Dániel Marx.